Deducing Truth from Correlation

J. Nötzel, W. Swetly

arXiv:1412.5831v4 [cs.IT] 4 Oct 2016

***Abstract*—This work is motivated by a question at the heart of unsupervised learning approaches:**
Assume we are collecting a number K of (subjective) opinions about some event E from K different agents. Can we infer E from them? Prima facie this seems impossible, since the agents may be lying.
We model this task by letting the events be distributed according to some distribution p and the task is to estimate p under unknown noise. Again, this is impossible without additional assumptions. We report here the finding of very natural such assumptions - the availability of multiple copies of the true data, each under independent and invertible (in the sense of matrices) noise, is already sufficient:
If the true distribution and the observations are modelled on the same finite alphabet, then the number of such copies needed to determine p to the highest possible precision is exactly three! This result can be seen as a counterpart to independent component analysis. Therefore, we call our approach 'dependent component analysis'.
In addition, we present generalizations of the model to different alphabet sizes at in- and output. A second result is found: the 'activation' of invertibility through multiple parallel uses.

Keywords: unsupervised learning, spatial diversity, blind source estimation, dependent component analysis, independent component analysis.

I. Introduction

Can we know the objective truth about the distribution of a set of events E?

This question can of course be cast into many different and more precise forms. We will be concerned here with the following version of it: We assume we are collecting a number K of potentially subjective opinions (e.g. from witnesses of some crime) about E from K different agents. A simple sketch of the scenario is given in the following figure.

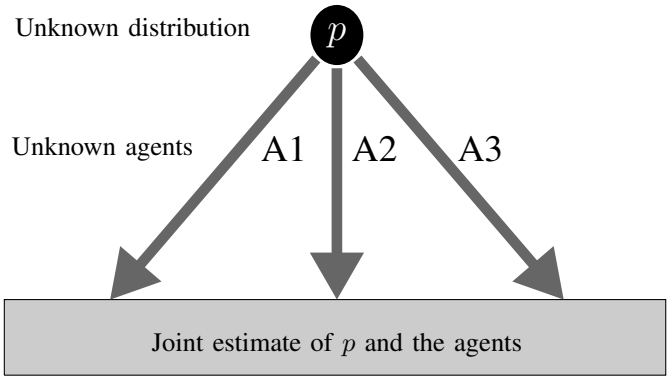

Fig. 1. Dependent component system with 3 agents

We show in this work that under certain conditions the task is then feasible: We can find out the objective truth from the subjective opinions given that these opinions are only minimally correlated to the true event, and that the different agents give their respective opinions independently from each other (they are not conspiring). It seems reasonable to assume that especially human perception is correlated to a common objective reality, if this objective reality exists. Thus, our model offers a new way of looking at the process in which a multitude of different opinions about the same objective truth can enable an observer having access to all these opinions to actually find the objective truth. We interpret our result as a mathematical statement in the favour of cooperative actions in the following sense: In the task of finding out an objective truth, one could follow the approach of finding an agent that reliably reports only true values, if necessary by building or training it first and then testing it in various situations. Our approach is contrary in nature: We do not seek to find such an agent at all. Instead, we build on the joint use of multiple, uncorrelated observations. In addition, no 'training' is necessary and no assumptions are made on the mutual information between in- and output of the channels. In fact, the mutual information between the true events and the events reported by the agents has to be nonzero but can be arbitrarily small otherwise.

From an information-theoretic perspective, which we adopt in this work, the task of finding the true state of some object or process is formulated best in terms of hypothesis testing. A huge amount of fundamental results has been

J. Nötzel had a joint affiliation with the Lehrstuhl für Theoretische Informationstechnik at Technische Universität München in 80290 München and Física Teòrica: Informació i Fenòmens Quàntics, Universitat Autònoma de Barcelona in ES-08193 Bellaterra (Barcelona), Spain. He is now with the Technische Universität Dresden, Communications Laboratory, 01069 Dresden, Germany.

W. Swetly is affiliated with the Munich Center for Mathematical Philosophy, Ludwig-Maximilians-Universität, 80539 München, Germany and has been with the Lehrstuhl für Datenverarbeitung, Technische Universität München, 80290 München, Germany during the preparation of the first draft of this work.

J.N. thanks Holger Boche for continuously supporting his research efforts, and W.S. thanks Stephan Hartmann and Klaus Diepold for the great opportunities from 2012 to 2014. He is especially grateful to his academic teachers Hannes Leitgeb, Godehard Link and Carlos Moulines for their continous and generous support. This work was supported by the BMBF (grant 01BQ1050, J.N.) and by the DFG (grants NO 1129/1-1 and BO 1734/20-1, J.N.) and the Cluster of Excellence CoTeSys (W.S.). Further funding (J.N.) was provided by the ERC Advanced Grant IRQUAT, the Spanish MINECO Project No. FIS2013-40627-P and the Generalitat de Catalunya CIRIT Project No. 2014 SGR 966. This paper was presented in part at the Internationl Congress on Mathematical Physics (ICMP) in Santiago de Chile in 2015.

©2016 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Digital Object Identifier 10.1109/TIT.2016.2611661

obtained e.g. in [9], [4], [10], [23]. A simple introduction to basic reasoning in this area can be found in [11].
However, the task of hypothesis testing is usually formulated such that a direct access to the source, system or process is guaranteed. In the light of developments e.g. in quantum theory or with an eye on extraterrestrial exploration, this seems highly questionable - the system to be observed is usually being observed via its interaction with a measurement apparatus, which in the event gets correlated to the system. The observer ultimately draws his conclusions from the output given to him through the measurement device.
One may now argue that the uncertainties in the measurement apparatus could in many situations be circumvented by adjusting it properly. This certainly requires that one makes measurements on already well-known inputs. One readily sees that this argument is circular in nature - how did one come to know these well-known inputs?
It seems reasonable to take one step back and consider the task of hypothesis testing under unknown noise. Such noise may for example be introduced through imperfections built in a measurement apparatus, or the lack of knowledge regarding events reported to us via not necessarily trusted agents.

The development of a complete theory of dependent components systems is way beyond the scope of this work. Although we aim at formulations in traditional hypothesis testing scenarios, there is one basic question which has to be answered at first and it is this question that we first pose and then solve here:

Is a dependent component system invertible?

Of course, if two different distributions p and p' of the events E could potentially get mapped to one and the same output distribution q by the influence of noise, any approach to differentiate between p and p' would be doomed to fail. Thus, it is of utmost importance to clarify this one point before starting to formulate more elaborate tasks.

Before we go into more detail, we now give a first and informal definition of the term 'dependent component analysis'. For simplicity, we will call the systems under consideration dependent component systems (DCS). Quite generally, such a system is to be understood as any physical system in a given state p, together with a number K of channels (linear, positivity-preserving maps going from the system to their respective output systems).
The goal of DCA is to determine, from data taken from all or some of the K channel outputs, the true state p of the system.
More specifically we will, throughout this work, assume that the system under consideration is given by a probability distribution p on a finite alphabet $\{1, \ldots, L\}$ which simply labels the events E (without loss of generality the events E are therefore given by natural numbers). Through the time of $n \in \mathbb{N}$ observations, the system generates the events $(E_1, \ldots, E_n)$ which are distributed independently and identically according to p. Each channel receives an exact copy of this sequence and transmits it to the output. The channels are assumed to be memoryless. They act independently from each other.
Our main result in this situation is the following: As long as the set of possible events at the output of each of the channels has the same cardinality as the set of possible events E at the input, the number K of channels satisfies $K \geq 3$ and each of them is invertible as a matrix, the distribution p of the events E can be inferred up to a permutation if one knows the distribution of events at the output (which can be approximated to arbitrary precision from observed data due to the assumed structure of the channels and distribution p). We additionally prove that even non-invertible channels can be used to obtain this result, if only enough of them are available.

Outline of the paper. We first state our main results in Section IV. These clarify when a DCA system can be inverted. We give examples for non-invertible systems as well. Thus an open question remains: Under which circumstances is a DCA system invertible, and can this be detected solely from observations at the output of the system *and* from knowing that it *is* in fact a DCA system?
A partial answer to this is given by our Theorem 4, which states that multiple parallel uses of one and the same channel can be inverted if the inputs are restricted to a certain form, even if the channel itself is non-invertible.
The proof of our statements are given in Section V. In the appendix (Section VI) we provide an additional subsection which highlights the connection to hypothesis testing and clarifies how the overall detection process can be carried out. This connects our approach to [4], [23], [10], [9] - our work is a first crucial step towards a generalization of hypothesis testing to situations where the test takes place under some additional, unknown noise. Finally in subsection VI-B we briefly connect to the Simpson-Yule paradox and the 'conjunctive fork'. Page numbers are as follows:

I	**Introduction**	1
II	**Notation, motivation and introduction of the framework of multilinear algebra**	4
III	**Definitions**	6
IV	**Main results, examples**	7
V	**Proofs**	9
VI	**Appendix**	12
VII	**Open problems**	12
References		13
Biographies		13

A. Historical notes and connections to other approaches

The reader interested in the subject will find a multitude of different approaches to systems with additional structure, like the one treated here. For example the approach taken in this work applies to multiple antenna systems and radar as well, and this is an area of research that already reported use of

the effect described here as early as 1931 in [2], [3], although the mathematical treatment given in this work differs strongly from these earlier approaches. Our approach is finally able to provide a deeper and very general understanding of the phenomenon from a clean perspective, if only at the price of a finite-dimensional analysis.

Another application that seems to fall into the category of dependent component systems is human perception: The different sensing systems (e.g. vision and hearing) can be assumed to be subjected to independent noise most of the time. Anything which affects both vision and hearing at the same time can, according to everyday-experience, in general be identified very precisely.

With multiple independent copies, perspectives and opinions on all kinds of subjects via the internet, dependent component analysis (which will usually be abbreviated by DCA in the following) can certainly be applied in data analysis as well, and at least in spirit this effect is exploited for noise estimation in digital images e.g. in the recent work [32].

As mentioned already, such studies date back at least as far as the 1930's. We therefore confine ourselves here to the mentioning of only a few research areas which we feel are important either from practical or theoretical, if not even philosophical perspectives. We also restrict ourselves to citing only very few published results in these areas, and we picked them such that the references contained therein enable the reader to quickly enter the corresponding field.

At first, let us mention the famous ICA (independent component analysis) approach, which can be considered orthogonal in spirit to ours. In ICA, the system under consideration consists of K independent parts, and the transformation between the system and the observer is only assumed to be linear.

The astonishing result in ICA is, that it is possible to detect both channel and system, up to a permutation, but only from observing the output and from knowing that the system under consideration fulfills above assumptions. For a good introduction to ICA, including its history, see [8] or [20].

Another branch which has to be mentioned here is the analysis of multichannel systems. An introduction to these topics can be found for example in [29] or [5]. A first paper summarizing different approaches to the topic was published by Brennan as early as 1959 [6]. Many contributions from the engineering perspective can be found under the keyword 'diversity combining'.

Surprisingly, it seems the situation has never been analyzed in an information-theoretic context. The results which are known to the authors consider several restoration problems, among them image restoration, but do not exploit the specific probabilistic structure itself nor do they consider the various problems arising from different alphabet sizes. Also, it seems to have slipped the attention of earlier research that dependent component systems can, under not too strong assumptions, be inverted. Thanks to the comments of an unknown reviewer, we were made aware of the publication [31] that treats the problem of denoising in an information-theoretic framework. In that work, a string $x^n = (x_1, \dots, x_n)$ representing undistorted information, is subjected to noise acting independently on every symbol x_i, producing the output $y^n = (y_1, \dots, y_n)$. An algorithm (called discrete universal denoiser or DUDE, for short) is presented that, when the parameters of the noise model are known, delivers optimal performance in the recovery of x^n from y^n. Of course those parameters may not always be given, in which case the algorithm is completely useless. However, if multiple copies of x^n have been stored in the past and then corrupted by independent noise, our algorithm is able to deliver exactly the noise model that serves as an input to the algorithm presented in [31]. The DUDE has been extended to cover situations involving channel uncertainty in [17].

Our model is further intimately connected to the study of Markov chains (see e.g. the work [14] for an overview), as our model can be viewed as a hidden "arbitrarily varying" Markov model in which the state space of the Markov chain equals its output space, but the information delivered to the observer is not via one and the same noisy channel, but via different channels:

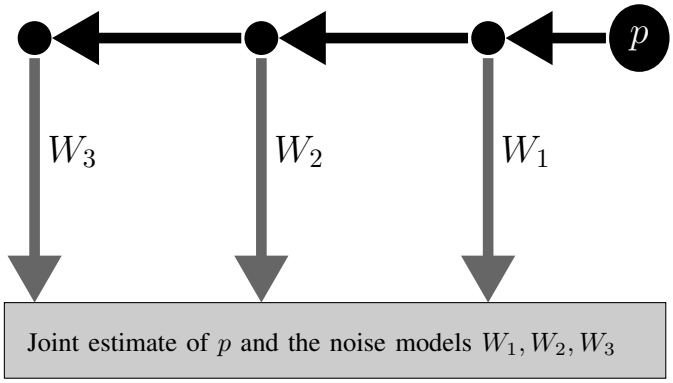

Fig. 2. Dependent component system seen as Markov chain with different (unknown) noisy channels W_1, W_2, W_3 and unknown input distribution p.

After finalization of the key results and statements of this manuscript we became aware that the name "dependent component analysis" has actually been used in the literature earlier already (see or example [33] or [24] and references therein for an overview), although not in the framework that is treated here. Rather, a multitude of different approaches dealing with estimation of multivariate distributions is presented there. To the author's knowledge, this work is the very first to pose the fundamental question of invertibility together with the question of how many independent observations one has to make in order to invert the system. That the answer to this question is exactly three is a surprising result, which we are at present tempted to not see as a mere artifact.

Invertibility of channel matrices has also played a minor role in the work [18] on approximation of output statistics of channels: After developing a theory of resolvability for output statistics of channels, the authors observed that statements about its output statistics translate to statements about the input statistics when the channel is invertible (see the discussion of [18, Theorem 15]). In a second paper [19], the authors of [18] explicitly exploited this observation to further develop the theory of resolvability by extending it to input processes

of a channel.
The type of question we study here is, at least in spirit, similar to the search for informationally complete measurements in quantum information theory [7], [12], [13], [26]. Such measurements have the property that they guarantee the possibility to distinguish between the many possible states that a physical system may be in.
From the recent work [25] which is inspired among others by results of C.F. von Weizsäcker [30] it is known that the quantum bit space (which can be represented as the unit ball in exactly three real dimensions) and the three-dimensional space structure we are experiencing every day can actually be related by a number of clearly specified reasonable assumptions and logical arguments. We hypothesize that similar arguments should make it possible to connect our findings to the geometry of space.

II. Notation, motivation and introduction of the framework of multilinear algebra

This section is split into three subsections. First we define the standard notation, a task which is rather repetitive. Then we give a more precise formulation of our problem statement, using the established notation. We explain which difficulties arise. Then, we introduce the framework of multilinear algebra which we use later to solve the problem.

A. Notation

For a natural number L, we define $[L] := \{1, \dots, L\}$. The set of permutations on $[L]$ is denoted S_L. Given two such sets $[L_1], [L_2]$, their product is $[L_1] \times [L_2] := \{(l_1, l_2) : l_1 \in [L_1],\ l_2 \in [L_2]\}$. For any natural numbers n and L, $[L]^n$ is the n-fold product of $[L]$ with itself. The set of real-valued functions on a finite set $[L]$ is $\mathfrak{F}_L$. The set of probability distributions on a finite set $[L]$ is

$$\mathfrak{P}([L]) := \{p \in \mathfrak{F}_L : p(i) \geq 0 \ \forall\ i \in [L],\ \sum_{i=1}^{L} p(i) = 1\}. \quad (1)$$

The support of $p \in \mathfrak{P}([L])$ is $\mathrm{supp}(p) := \{i \in [L] : p(i) > 0\}$. A very important subset of elements of $\mathfrak{P}([L])$ is the set of its extreme points, the Dirac-measures: for $i \in [L]$, $\delta_i \in \mathfrak{P}([L])$ is defined through $\delta_i(j) = \delta(i, j)$, where $\delta(\cdot, \cdot)$ is the usual Kronecker-delta symbol. Closely related to the latter is the set of indicator functions on $[L]$. These are defined for subsets $S \subset [L]$ via $\mathbb{1}_S(i) = 1$ if and only if $i \in S$ and $\mathbb{1}_S(i) = 0$ else.
Two subsets of $\mathfrak{P}([L])$ which are of great importance in our case are

$$\mathfrak{P}_>([L]) := \{p \in \mathfrak{P}([L]) : p(i) > 0 \ \forall\ i \in [L]\} \quad (2)$$

$$\mathfrak{P}^{\downarrow}([L]) := \{p \in \mathfrak{P}([L]) : p(1) > \ldots > p(L)\}. \quad (3)$$

The main task that we will be dealing with is that of *distinguishing* between elements $p, p' \in \mathfrak{P}([L])$. While there are many ways of doing this, we would like to single out two of them here: First, we may use (in principle for any number $k \in \mathbb{R}$, though we will make use only of $k = 1$ and $k = 2$ here) the k-norms $\|p - p'\|_k := (\sum_{i=1}^{L} |p(i) - p'(i)|^k)^{1/k}$. If $k = 1$, we will omit the subscript in the following.
Second, one may use the relative entropy or Kullback-Leibler distance: $D(p\|p')$ is defined as $D(p\|p') := \sum_{i=1}^{L} p(i) \log(p(i)/p'(i))$ if $p'(i) = 0 \Rightarrow p(i) = 0\ \forall i \in [L]$ and $D(p\|p') := +\infty$ if this is not the case.
The noise that complicates the task of estimating p is modelled by matrices W of conditional probability distributions $(w(i|j))_{i,j=1}^{L',L}$ whose entries are numbers in the interval $[0, 1]$ satisfying, for all $j \in [L]$, $\sum_{i=1}^{L} w(i|j) = 1$. The set of all such matrices is denoted $\mathcal{W}([L], [L])$. Such matrices are, using standard terminology of Shannon information theory, equivalently called a "channel".
We will later relate these probabilistic concepts to linear algebra in a more explicit fashion. In order to motivate this approach, we first restate the problem using the rather generic language that we introduced so far:

B. Problem statement (standard formalism)

Let $L, K \in \mathbb{N}$ be given. To a given choice $W_1, \dots, W_K \in \mathcal{W}([L], [L])$ of channels and $p \in \mathfrak{P}_>([L])$ we let $q \in \mathfrak{P}([L]^K)$ be the distribution defined by setting, for all $y_1, \dots, y_K \in [L]$,

$$q(y_1, \dots, y_K) = \sum_{x \in [L]} p(x) \prod_{k=1}^{K} w_k(y_k|x). \quad (4)$$

The question we ask is, for what other choices $p' \in \mathfrak{P}([L])$ and $W_1, \dots, W_k \in \mathcal{W}([L], [L])$ the equalities

$$q(y_1, \dots, y_K) = \sum_{x \in [L]} p'(x) \prod_{k=1}^{K} w'_k(y_k|x) \quad (5)$$

can hold for all $y_1, \dots, y_k \in [L]$. It is clear that the answer is completely unsatisfying for $K = 1$: For example the choice of the distribution p' satisfying (5) is completely independent from that of p, since it is always possible to choose an appropriate W'_1 such that (5) is fulfilled. Later in example 1 we explain shortly why the answer for $K = 2$ turns out to be not very satisfying as well. With the help of our Theorem 1 however we can demonstrate that, starting from $K = 3$, the solutions to the L^K equations (one for each choice of $y_1, \dots, y_K \in [L]$):

$$\sum_{x \in [L]} p(x) \prod_{k=1}^{K} w_k(y_k|x) = \sum_{x \in [L]} p'(x) \prod_{k=1}^{K} w'_k(y_k|x) \quad (6)$$

can be written as follows: Take any choice $p, W_1, \dots, W_K$. For every permutation $\tau : [L] \to [L]$, the distribution p' defined by $p'(i) := p(\tau(i))$ and channels W'_k defined by $w'_k(j|i) := w_k(j|\tau^{-1}(i))$ are a solution to the equation (6), and for a fixed choice of p and $W_1, \dots, W_K$ these are the only solutions.
Proving this result turned out to be a challenging task: Starting with restricted models where $L = 2$ and all channels W_k are restricted to lie in the class of binary symmetric channels $W_k = BSC_{s_k}$ with unknown probability s_k of causing a bit to be flipped, we realized that the polynomial equation in the variables $p(1), s_1, s_2$, and $p'(1), s'_1, s'_2$ arising from (6) has only finitely many solutions. We were also able to picture the manifolds arising from a fixed choice of $p(0)$ and variations of

s_1, s_2 and realized that they were forming a two-dimensional manifold.

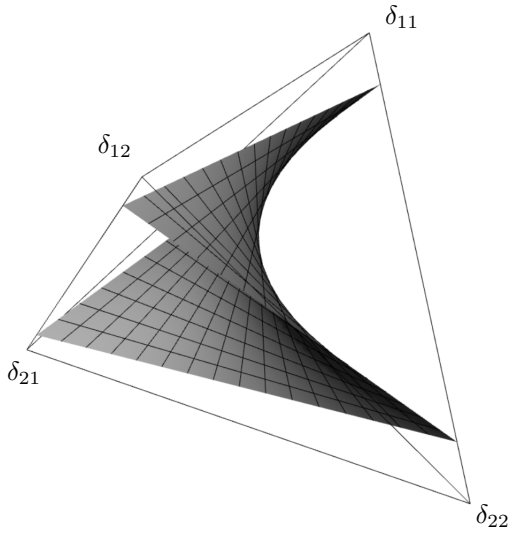

Fig. 3. Probability simplex over the alphabet $\{11, 12, 211, 22\}$. The δ_{ij} denote the Dirac-distributions corresponding to the respective letters in the alphabet. The two-dimensional manifold arises from setting $p(1) = 0.12$ and $K = 2$ in (6) and letting the channels w_1 and w_2 be binary-symmetric channels with flip probabilities $(s_1, s_2) \in [0, 1] \times [0, 1]$.

This motivated further study of the question. Unfortunately we were not able to solve the resulting polynomial equations in the many variables that arise as soon as the model becomes even slightly more complex. From the above picture it became clear that $K = 2$ would not be sufficient for more complex noise models. In fact, counting parameters led us to conclude the following: As any channel W introduces $(L-1)\cdot L$ variables, and p another $L-1$ variables, we could only hope to gain insightful solutions whenever the total number $K\cdot(L-1)\cdot L+L-1$ of variables was smaller than or equal to the dimension of the convex set $\mathfrak{P}([L])^K$, which equals $L^K - 1$. For natural numbers L, K, the inequality

$$L^K \geq K \cdot (L-1) \cdot L + L \tag{7}$$

however is easily seen to hold true for all $K \geq 3$. Unfortunately, already the binary case $L = 2$ with arbitrary channels $W_1, W_2, W_3, W_1', W_2', W_3' \in \mathcal{W}(2, 2)$ introduces an untractable total number of $2 \cdot 6 + 2 = 14$ variables, and no results concerning polynomials of the specific form introduced by our model could be found. After numerous attempts we finally reformulated the problem using the formalism that was developed during the evolution of quantum information theory over the last 20 years, and this approach finally turned out to be fruitful and helped us to reduce the problem to a small number of rather simple observations. The necessary notation for the reformulation of the problem is given in the following subsection.

C. Multilinear framework

We will from now on consider any $\mathfrak{P}([L])$ as being embedded into $\mathbb{R}^L$ through the bijective map $\Phi : \mathfrak{P}([L]) \to \mathbb{R}^L$ defined through $\Phi(\delta_i) := e_i$, where $\{e_i\}_{i=1}^L$ is any fixed orthonormal basis of $\mathbb{R}^L$. The image $\Phi(\mathfrak{P}([L]))$ naturally generates its $L-1$-dimensional supporting hyperplane within $\mathbb{R}^L$.

This embedding allows a natural use of matrix calculus. Since we will be dealing with composite systems throughout, an important part of our work requires basic results from multilinear algebra. We will now introduce these for bipartite systems, the generalization to the multipartite case is straightforward.

Throughout, we use one fixed basis $\{e_i\}_{i=1}^L$ for $\mathbb{R}^L$. L times L' matrices are thought of as linear maps from $\mathbb{R}^L$ to $\mathbb{R}^{L'}$ via their action in this basis. The set of linear maps from $\mathbb{R}^L$ to $\mathbb{R}^{L'}$ (matrices with L' rows and L columns) is denoted $M(L, L')$, the group of invertible matrices (if $L = L'$) is written $Gl(L)$. The range of a matrix $M \in M(L, L')$ is $\operatorname{ran} M := \{Mx : x \in \mathbb{R}^L\}$, and its kernel $\ker M$ (as usual) the set $\{x \in \mathbb{R}^L : Mx = 0\}$.

The scalar product $\langle \cdot, \cdot \rangle$ on $\mathbb{R}^L \times \mathbb{R}^L$ is the standard one: $\langle e_i, e_j \rangle = \delta(i, j)$. If needed, we will represent an $L' \times L$ matrix by using the matrix basis $\{E_{i,j}\}_{i,j=1}^{L',L}$ which is defined by $E_{i,j}e_k = \delta(k, j)e_i \ \forall i \in L',\ k, j \in L$. The identity matrix on $\mathbb{R}^L$ (which is at the same time a channel from $[L]$ to $[L]$) is written as $\mathbb{1} = \sum_{i=1}^L E_{ii}$. Care will be taken that it does not get confused with indicator functions (symbols of the form $\mathbb{1}_S$) that we defined earlier.

Without going into any further detail, we introduce the tensor product of $\mathbb{R}^L$ with $\mathbb{R}^K$ in a very straightforward manner by setting

$$\mathbb{R}^L \otimes \mathbb{R}^K := \operatorname{span}\{e_i \otimes e_k\}_{i=1,k=1}^{L,K}. \tag{8}$$

This allows us to define general 'product vectors' of two vectors $u = \sum_{i=1}^L u_i e_i$ and $v = \sum_{i=1}^K v_i e_i$ by $u \otimes v := \sum_{i,j=1}^{L,K} u_i v_j e_i \otimes e_j$. The vector space $\mathbb{R}^L \otimes \mathbb{R}^K$ inherits the scalar products from $\mathbb{R}^L$ and $\mathbb{R}^{L'}$ by the formula $\langle u \otimes v, x \otimes y \rangle := \langle u, x \rangle \langle v, y \rangle$. The corresponding matrix spaces are denoted $M((L, K), (L', K'))$. Given $A \in M(L, L')$ and $B \in M(K, K')$, we define $A \otimes B \in M((K, L), (K', L'))$ through its action on product vectors:

$$(A \otimes B)(u \otimes v) := (Au) \otimes (Bv). \tag{9}$$

This leads to the following set of rules:

$$\forall A, B \in M(L),\ C, D \in M(L) : \\ (A + B) \otimes (C + D) = \\ A \otimes C + A \otimes D + B \otimes C + B \otimes D, \tag{10}$$

$$\forall A \in Gl(K),\ B \in Gl(L) : \\ (A \otimes B)^{-1} = A^{-1} \otimes B^{-1} \tag{11}$$

$$\forall A \in M(K),\ B \in M(L),\ \alpha \in \mathbb{C} : \\ \alpha(A \otimes B) = (\alpha A) \otimes B = A \otimes (\alpha B) \tag{12}$$

$$\mathbb{1}_{\mathbb{R}^K \otimes \mathbb{R}^L} = \mathbb{1}_{\mathbb{R}^K} \otimes \mathbb{1}_{\mathbb{R}^L}. \tag{13}$$

In order to simplify notation later we will use, for $u \in \mathbb{R}^L$ and $n \in \mathbb{N}$, the shorthand $u^{\otimes n} := u \otimes \ldots \otimes u$ for the n-fold tensor product of u with itself. Accordingly, for $A \in M(L, K)$ we write $A^{\otimes n}$ to denote the n-fold product $A \otimes \ldots \otimes A$. A very important object we shall encounter is the vector $v_+ := \sum_{i=1}^L e_i \otimes e_i \in \mathbb{R}^L \otimes \mathbb{R}^L$. It has the important property that

$$(A \otimes \mathbb{1})v_+ = (\mathbb{1} \otimes A^\top)v_+ \qquad \forall A \in M(L), \tag{14}$$

where the matrix transposition $A \mapsto A^\top$ is defined in the basis $\{e_i\}_{i=1}^L$.

An important operation on a composite system described by the set of probability distributions $\mathfrak{P}([K] \times [L])$ is the partial trace $\mathrm{tr}_1 : \mathbb{R}^K \otimes \mathbb{R}^L \to \mathbb{R}^L$ which 'forgets' the information stored in the first system modelled over $\mathbb{R}^K$ by summing over it: For $a = \sum_{i,j=1}^{K,L} a_{i,j} e_i \otimes e_j$, it is defined by

$$\mathrm{tr}_1(a) := \sum_{i,j=1}^{K,L} a_{i,j} e_j. \tag{15}$$

This operation has the following nice property: If $v = \sum_{i,j=1}^{K,L} p(i,j) e_i \otimes e_j$ for some $p \in \mathfrak{P}([K] \times [L])$, then

$$\mathrm{tr}_1(v) = \sum_{i=1}^{L} \left(\sum_{j=1}^{K} p(i,j) \right) e_i, \tag{16}$$

which is clearly just the image of one of the marginal distributions of p under the (injective) map Φ. The notion easily extends to multi-partite systems $\mathbb{R}^{L_1} \otimes \ldots \otimes \mathbb{R}^{L_m}$ and defines tr_i for every $i \in [m]$ and, more generally, tr_A for every $A \in [m]$. Note that if only one system being modelled by $p \in \mathfrak{P}([L])$ is present we can as well abbreviate tr_1 by tr and then it holds $\mathrm{tr}(p) = \sum_{i=1}^L p(i) = 1$.

This way we have come back to probability distributions and are now in the position to embed the most natural linear maps on them into our framework: channels. A channel is a positivity and trace preserving linear map $W : \mathfrak{P}([L]) \to \mathfrak{P}([L'])$, where $L, L' \in \mathbb{N}$ are arbitrary. We may think of W as a matrix defined by its entries $w_{ij} := W(\delta_j)(i)$. If necessary and unambiguous, we shall also write $w(i|j) := w_{ij}$.

An important example of a channel is a permutation matrix $\tau \in S_L$: Its natural action as $\tau \in \mathcal{W}([L],[L])$ is via $\tau(p)(i) := p(\tau^{-1}(i))$.

Clearly, every channel is completely represented by the matrix with entries w_{ij} and application of W to a probability distribution $p \in \mathfrak{P}([L])$ is equivalent to applying the matrix defined via its matrix entries $w_{ij} := w(i|j)$ to the vector $\sum_{i=1}^L p(i) e_i$.

During our proofs, we will not necessarily always be working with channels, but rather with matrices. We will therefore spend a few more words on this connection.

It is clear that $\mathfrak{P}([L]) \subset \mathbb{R}_1^L := \{\sum_{i=1}^L v_i e_i \in \mathbb{R}^L : \sum_{i=1}^L v_i = 1\}$. Therefore, we have the implication $W \in \mathcal{W}([L],[L]) \Rightarrow W(\mathbb{R}_1^L) \subset \mathbb{R}_1^L$, and in case that W is invertible this lets us conclude that $W^{-1}(\mathbb{R}_1^L) = \mathbb{R}_1^L$.

We see that, as long as we restrict our analysis to matrices which are composed of channels or inverses of channels, all we need to take care of is their action on $\mathbb{R}_1^L$.

In addition, a channel $W \in \mathcal{W}([L],[L])$ is invertible if and only if the corresponding matrix $W \in M(L,L)$ is invertible. While it is clear that a non-invertible channel has a non-invertible matrix associated to it, the other direction of the claim can be established as follows:

Let for some real numbers $\gamma_1, \ldots, \gamma_L$, not all of which are zero, $\sum_{i=1}^L \gamma_i W(\delta_i) = 0$. Set $g := \sum_{i=1}^L \gamma_i$. If $g \neq 0$ we may set $\bar{\gamma}_i := \gamma_i / g$ for each $i \in [L]$ and conclude that $0 = \sum_{i=1}^L \bar{\gamma}_i W(\delta_i)$, in contradiction to $W(\mathbb{R}_1^L) \subset \mathbb{R}_1^L$. If $g = 0$ we define the number $\gamma_{\max} := \max_{i \in [L]} |\gamma_i|$. For every $\alpha \in \mathbb{R}$, we introduce the vector $\gamma_\alpha := \sum_{i=1}^L (\alpha \cdot \gamma_i + L^{-1}) \delta_i$. For each α it is clear that $\gamma_\alpha \in \mathbb{R}_1^L$. We additionally have

$$(\gamma_\alpha)_i = \alpha \cdot \gamma_i + L^{-1} \tag{17}$$

$$\geq L^{-1} - |\alpha| \cdot \gamma_{\max} \tag{18}$$

$$\geq 0 \tag{19}$$

whenever $\alpha \in [0, \frac{1}{L \cdot \gamma_{\max}}]$, meaning that $\gamma_\alpha \in \mathfrak{P}([L])$ for that range. By assumption $W(\alpha \cdot \gamma) = 0$ for all α and this implies that especially $W(\gamma_\alpha) = W(\gamma_0)$ for all $\alpha \in [0, \frac{1}{L \cdot \gamma_{\max}}]$. But at least one of the γ_i is nonzero, thus $\gamma_\alpha \neq \gamma_{\alpha'}$ for all $\alpha, \alpha' \in [0, \frac{1}{L \cdot \gamma_{\max}}]$ for which $\alpha \neq \alpha'$. This is in contradiction to the assumed invertibility of W as a map from $\mathfrak{P}([L])$ to $\mathfrak{P}([L])$.

III. Definitions

We are now able to formalize our problem using the language of multilinear algebra. Throughout our discussion, we will assume the existence of a probability distribution p on some finite set $[L]$ that models the 'true' state of a physical system which emits signals $i \in [L]$ independently and identically distributed according to p.

Since p is fixed, it makes little sense to talk about events which do not happen according to p. Therefore, we will always assume that $p(i) > 0$ holds for all $i \in [L]$. In other words: $p \in \mathfrak{P}_>([L])$.

On the other hand, this requires us to 'learn' the parameter L of the system as well. We therefore list three different scenarios: We start with the case $L = L'$, which is central to the whole discussion and delivers the necessary tools for a discussion of the other cases, namely $L' < L$ and $L' > L$.

We will now list the definitions that we need in order to state our results. First, a technical thing:

Definition 1. *Given $p \in \mathfrak{P}([L])$ we denote by $p^{(K)}$ the distribution*

$$p^{(K)} := \sum_{i=1}^{L} p(i) \delta_i^{\otimes K}. \tag{20}$$

For any fixed $[L]$, the set of all such $p^{(K)}$ is

$$\mathfrak{P}^{(K)}([L]) := \{p^{(K)} : p \in \mathfrak{P}([L])\}. \tag{21}$$

The main definition is the following.

Definition 2 (Dependent component system (DCS))**.** *For given natural numbers L, L' and K, we define For given natural numbers L, L' and K, we define $DCS(L,K,L')$ to be the set of all $(p, W_1, \ldots, W_K)$ such that $p \in \mathfrak{P}([L])$ and $W_1, \ldots, W_K \in \mathcal{W}([L],[L'])$. This is the set of all dependent component systems. Any element of it is represented (uniquely only if $p \in \mathfrak{P}_>([L])$) by the distribution*

$$\left(\mathbb{1} \otimes \bigotimes_{i=1}^{K} W_i \right) p^{(K+1)} \in \mathfrak{P}([L] \times [L']^K). \tag{22}$$

In order to shorten notation we will usually write sentences like 'let a $DCS(L,K,L')$ be given', implying that the mathematical object under study is a system $\mathcal{S} \in DCS(L,K,L')$. For a more thorough analysis of dependent component systems we need additional definitions:

Definition 3. *For $L, L', K \in \mathbb{N}$ we define the $DCS(L,K,L')$ surface to be*

$$\{(\otimes_{i=1}^K W_i)p^{(K)} : (p, \mathbf{W}) \in DCS(L,K,L')\}, \tag{23}$$

where $(p, \mathbf{W}) = (p, W_1, \ldots, W_K)$. As an important subset of this surface we consider the $FR(L,K,L')$ of all distribtions $(\otimes_{i=1}^K W_i)p^{(K)} \in \mathfrak{P}([L']^K)$ such that $p \in \mathfrak{P}_>([L])$ and $\mathrm{ran}(W_1) = \ldots = \mathrm{ran}(W_K) = \mathbb{R}^{L'}$. This is the set of those points which are generated by full-range channels and strictly positive distributions.
In case $L' \leq L$ we are going to need the subset $FRSK(L,K,L')$ of all distributions $(\otimes_{i=1}^K W_i)p^{(K)} \in \mathfrak{P}([L']^K)$ such that $p \in \mathfrak{P}_>([L])$ and both $\mathrm{ran}(W_i) = \mathbb{R}^{L'}$ and $\ker(W_i) = \ker(W_j)$ for all $i, j \in [K]$.
This subset of $FR(L,K,L')$ consists of all those points on the $DCS(L,K,L')$ surface which are generated by strictly positive probability distributions and K full range channels such that all of them have the same kernel.

Remark 1. *Of course, $FR(L,K,L') = \emptyset$ for $L' > L$.*

Let us have a short look at an insightful example.

Example 1 ($FR(L,2,L)$)**.** *Let $q \in \mathfrak{P}([L]^2)$. It can obviously be written as $q(i,j) = p(i)r(j|i)$. If there is no such decomposition such that $\mathrm{supp}(p) = [L]$ and r is invertible, then we may choose an $\varepsilon > 0$ as small as we like and an invertible r' together with $\mathrm{supp}(p') = [L]$ such that q' defined via $q'(i,j) := p'(i)r'(j|i)$ satisfies $\|q' - q\| \leq \varepsilon$. Take (p', Id, r') as a $DCS(L,2,L)$ system. Then*

$$(r' \otimes Id)p'^{(2)} = \sum_{i=1}^{l} p'(i) \sum_{j,k=1}^{L} r'(j|i)\delta(k,i)\delta_j \otimes \delta_k \tag{24}$$

$$= \sum_{i,j=1}^{L} r'(j|i)p'(i)\delta_j \otimes \delta_i \tag{25}$$

$$= q'. \tag{26}$$

It follows that $FR(L,2,L)$ is dense in $\mathfrak{P}([L]^2)$. The same holds true for $K = 1$.

This example clearly demonstrates that practically every distribution in $\mathfrak{P}([L]^2)$ can arise as an output of a dependent component system. This serves as another motivation to look for solutions to our problem statement (6) only for $K \geq 3$, thus affirming the intuition that one gets from parameter counting as in (7).

IV. Main results, examples

Our main result is the following Theorem 1, which has to be read with the following in mind: If $W_1, \ldots, W_K, V_1, \ldots, V_K \in \mathcal{W}([L],[L])$ are all invertible as matrices and $X_i := V_i^{-1} \circ W_i$ $(i = 1, \ldots, K)$, then the matrices $X_i \in M(L,L)$ are invertible and they map $\mathbb{R}_1^L$ to $\mathbb{R}_1^L$. Let the set of all matrices which map $\mathbb{R}_1^L$ to itself be $Gl_1(L)$. Another way of writing this is to set $Gl_1(L) := \{X : \sum_{j=1}^L X_{ji} = 1 \ \forall \ i \in [L]\}$. The exact connection between Theorem 1 and our initial question (6) is explained in detail at the end of the proof of Theorem 1.

Theorem 1. *[Uniqueness of Solution for $L = L'$] Let $K \in \mathbb{N}$ satisfy $K \geq 3$. Let $p \in \mathfrak{P}_>([L])$. There are exactly $L!$ tuples $(X_1, \ldots, X_L, p')$ of matrices $X_1, \ldots, X_K \in Gl_1(L)$ and probability distributions $p' \in \mathfrak{P}_>([L])$ satisfying the equation*

$$\sum_{i=1}^{L} p(i)\delta_i^{\otimes K} = (X_1 \otimes \ldots \otimes X_K)\left(\sum_{i=1}^{L} p'(i)\delta_i^{\otimes K}\right). \tag{27}$$

These are as follows: For every $\tau \in S_L$, the matrices $X_1 = \ldots = X_K = \tau^{-1}$ and $p' = \tau(p)$ solve (27), and these are the only solutions.
As a consequence, the function $\Theta : \mathfrak{P}([L]) \times \mathcal{W}([L],[L])^K \to \mathfrak{P}([L]^K)$ defined by

$$\Theta[(p, T_1, \ldots, T_K)] := \left(\bigotimes_{i=1}^{K} T_i\right)\left(\sum_{i=1}^{L} p(i)\delta_i^{\otimes K}\right) \tag{28}$$

is invertible if restricted to the right subset: there exists a function $\Theta' : DCS(L,K,L) \to \mathfrak{P}^{\downarrow}([L])$ that has the property

$$\Theta'(\Theta[(p, T_1, \ldots, T_K)]) = p, \tag{29}$$

for all $p \in \mathfrak{P}^{\downarrow}([L]) \cap \mathfrak{P}_>([L])$ and those $(T_1, \ldots, T_K)$ for which every T_i, $i \in [K]$, is invertible. In addition to that, there exists a second function $\Theta'' : DCS(L,K,L) \to \mathfrak{P}^{\downarrow}([L]) \times \mathcal{W}([L],[L])^K$ such that

$$\Theta''(\Theta[(p, T_1, \ldots, T_K)]) = (p, T_1, \ldots, T_K) \tag{30}$$

for all $p \in \mathfrak{P}^{\downarrow}([L]) \cap \mathfrak{P}_>([L])$ and those $(T_1, \ldots, T_K)$ for which every T_i, $i \in [K]$, is invertible.

Remark 2. *The following Theorem 2 shows that Θ' has the inversion property (29) even for all $p \in \mathfrak{P}^{\downarrow}([L])$ and those $T_1, \ldots, T_K$ that are invertible if restricted to $\mathrm{span}(\{e_i : p(i) > 0\})$.*
We will give a proof of this theorem for the interesting case $K = 3$ only, the general case offers no increase in insight.
It will become apparent from the proof that the theorem can be extended to the case where p, p' are not necessarily probability distributions (in which case the X_i will not necessarily be permutations any more).

In addition, we provide a generalization to the case where the input alphabet is strictly larger than the output alphabet. This situation should be considered the generic case in all sensor networks that involve a digital-analog converter.

Theorem 2 (The perfect conspiracy)**.** *Let $L > L'$ and $K \geq 3$. Then*

$$FR(L',K,L') \cap FR(L,K,L') = FRSK(L,K,L'). \tag{31}$$

Remark 3. *The implication of the theorem is the following important 'rule of thumb': If you observe three outputs of dimension L', and you can verify that your observed distribution q is in $FR(L',K,L')$ then either $L = L'$ or there is a 'perfect conspiracy' between all the channels in the sense that they all delete the same information.*
If you are happy that there are no apparent contradictions in your system, you may just leave it the way it is and conclude that the truth is given by some p', which may e.g. be given by a projection of the true p onto some hyperplane within $\mathfrak{P}([L])$.

If on the contrary you are a wary individual, you may always add new channels to the system, suspecting that in fact $L' > L$ and that it will be possible to find a channel which does not participate in the perfect conspiracy.

We now consider the case $L' > L$, which can also be interpreted as generalization of Theorem 1 to the singular cases (e.g. those cases where $p \notin \mathfrak{P}_>([L])$).

Theorem 3 (The case $L' > L$)**.** *Any $DCA(L, K, L')$ system with at least three channels and $L' > L$ is invertible up to a permutation on $[L]$.*

All the previous results are exact, and yet leave us unsatisfied: Assume we observe a set of K outputs of channels with output alphabet $[L']$, restrict our attention e.g. to all the triples of 3 subsystems and from Theorem 2 we infer that $L' < L$ has to hold (because the observed outputs may sometimes not be in $FR(L', K, L')$).
Then how can we use this knowledge to calculate p?
More specifically, is there any hope that a system of K channels $W_1, \ldots, W_K \in \mathcal{W}([L], [L'])$ can be invertible although $L' < L$ holds? Although we are not yet able to give a complete solution to this question including dependencies of the parameters K, L, L' and potentially necessary assumptions concerning p and the channels $W_1, \ldots, W_K$, we can already answer it in the affirmative.
To this end, consider a specific example: A channel $W \in \mathcal{W}([3], [2])$. The generic situation we encounter in this case will be that

$$W(\delta_3) = \lambda W(\delta_1) + (1-\lambda)W(\delta_2), \quad W(\delta_1) \neq W(\delta_2) \tag{32}$$

for some $\lambda \in [0, 1]$ and up to a permuation on $[3]$. We would like to find out now whether the three vectors $\{W(\delta_i) \otimes W(\delta_i)\}_{i=1}^3$ form a linearly independent set. If that is so, their supporting hyperplane has dimension two, just like $\mathfrak{P}([3])$. Hence, $W \otimes W$ would be invertible as a map from $\mathfrak{P}^{(2)}([3])$ to its image in $\mathfrak{P}([2]^2)$.
Assume that, on the contrary, there are $\gamma_1, \gamma_2, \gamma_3 \in \mathbb{R}$ not all of which are zero and such that

$$0 = \sum_{i=1}^{3} \gamma_i W(\delta_i) \otimes W(\delta_i) \tag{33}$$

holds. This is (introducing the abbreviations $w_1 := W(\delta_1),\ w_2 := W(\delta_2)$) equivalent to

$$\begin{aligned} 0 = (\gamma_1 + \lambda^2\gamma_3)w_1^{\otimes 2} + (\gamma_2 + (1-\lambda)^2\gamma_3)w_2^{\otimes 2} \\ + \gamma_3\lambda(1-\lambda)(w_1 \otimes w_2 + w_2 \otimes w_1). \end{aligned} \tag{34}$$

In case that $\lambda \notin \{0, 1\}$ this implies $\gamma_3 = 0$, which then leads to $\gamma_1 = \gamma_2 = 0$ as well and the desired result is proven by contradiction.
If, however, $\lambda \in \{0, 1\}$ holds then the above argument does not work. Then, it is even true that $W \otimes W$ is not invertible, since (w.l.o.g. $\lambda = 0$)

$$\begin{aligned} W \otimes W(p^{(2)}) = p(1)W(\delta_1) \otimes W(\delta_1) \\ + (p(2) + p(3))W(\delta_2) \otimes W(\delta_2). \end{aligned} \tag{35}$$

Only in that case is every information about the difference between 2 and 3 destroyed at the output and impossible to recover!
Let us now become slightly more general. Consider $W \in \mathcal{W}([L], [L'])$ and assume that the $W(\delta_i)$, $i = 1, \ldots, L'$, are pairwise different. We ask, when exactly can $W^{\otimes K}$, restricted to $\mathfrak{P}^{(K)}([L])$ be invertible? A partial answer is the following theorem:

Theorem 4. *Let $W \in \mathcal{W}([L], [L'])$ satisfy $W(\delta_i) \neq W(\delta_j)$ for all $i \neq j \in [L]$. Then $K \geq L - 1$ is sufficient for $W^{\otimes K}$ to be invertible as a map from $\mathfrak{P}^{(K)}([L])$ to $\mathfrak{P}([L'])$.*

We would like to express our belief that this is at the same time already optimal through the following conjecture:

Conjecture 1. *Under the preliminaries of Theorem 4 we have the following: If $K < L - 1$, then there is always a channel $W \in \mathcal{W}([L], [L'])$ satisfying $W(\delta_i) \neq W(\delta_j)$ for all $i \neq j \in [L]$ and such that $W^{\otimes K}$ is not invertible as a map from $\mathfrak{P}^{(K)}([L])$ to $\mathfrak{P}([L'])$.*

How do we solve the problem of inverting a DCA system concretely? Given that we estimated $\hat{q} \in \mathfrak{P}([L']^K)$, an easy solution which can be computed efficiently is the convex optimization problem

$$\arg \min_{W_1, \ldots, W_K, p} D((\otimes_{i=1}^K W_i)p^{(K)} \| \hat{q}). \tag{36}$$

If $\hat{q}$ is indeed an output distribution of a DCA system, above algorithm will return the system (up to a permutation). The ambiguities in the solution can be reduced by optimizing not over all $p \in \mathfrak{P}([L])$, but only over $\mathfrak{P}^{\downarrow}([L])$.
Of course, it is not at all necessary to use Kullback-Leibler divergence in the above, one could as well use any norm $\|\cdot\|$ and instead compute

$$\arg \min_{W_1, \ldots, W_K, p} \|(\otimes_{i=1}^K W_i)p^{(K)} - \hat{q}\|. \tag{37}$$

For practical purposes, it may even be useful to use a smooth quantity like $\|\cdot\|_2^2$ defined by $\|x\|_2 := \sqrt{\sum_{i=1}^d x_i^2}$, where $d = K \cdot L$ in our application and we employ the embedding ϕ of $\mathfrak{P}([L]^K)$ into $\mathbb{R}^{K \cdot L}$ introduced in Section II.

It may be speculated whether there is a simpler description of the $DCS(L, K, L')$ than the one given through calculation of all its points $(W_1 \otimes \ldots W_K)p^{(K)}$. Especially mutual information has turned out to be an important concept in various scenarios. Since the original source p can, in our model, only be accessed via the channel $(\otimes_{i=1}^K W_i)$, it makes little sense to look at mutual information between the in- and the output of the system. We provide here a statement showing that the pairwise mutual informations of the output systems are also not relevant quantities in our scenario:

Lemma 1 (Positivity of Mutual Information is not Sufficient)**.** *Consider $L = 4$, and let $L' \geq 2$. For every $K \in \mathbb{N}$, there exist channels $W_1, \ldots W_K \in \mathcal{W}([L], [L'])$ such that the random variables $Y_1, \ldots, Y_K$ defined by $\mathbb{P}(Y_1, \ldots, Y_K = i_1, \ldots, i_K) := (\otimes_{i=1}^K W_i)p^{(K)}(i_1, \ldots, i_K)$ satisfy $I(Y_i; Y_j) > 0$ for all $i \neq j \in \{1, \ldots, K\}$ and $p \in \mathfrak{P}_>([L])$, but p cannot be inferred from $(Y_1, \ldots, Y_K)$.*

V. Proofs

This section contains the proofs of our results, in the same order as they were stated in the previous section.

Proof of Theorem 1: As mentioned already, we restrict the proof to the case $K = 3$. The general case is a straightforward generalization. We start by considering the bivariate cases:

$$p^{(2)} = X_i \otimes X_j(p'^{(2)}), \qquad i \neq j. \tag{38}$$

These arise from the case with $K = 3$ by taking the partial trace. In this case, since all the alphabets are equal, tr_i denotes the trace over the i-th copy of $[L]$. For example for $i = 1$, $j = 2$ and $k = 3$ this works as follows: First we have

$$\begin{aligned} p^{(2)} &= \sum_{i=1}^{L} p(i)\delta_i \otimes \delta_i && (39)\\ &= \sum_{i=1}^{L} p(i)\mathrm{tr}_3\delta_i \otimes \delta_i \otimes \delta_i && (40)\\ &= \mathrm{tr}_3 \sum_{i=1}^{L} p(i)\delta_i \otimes \delta_i \otimes \delta_i && (41)\\ &= \mathrm{tr}_3 p^{(3)} && (42) \end{aligned}$$

and, since for every $i \in [L]$ we have

$$\mathrm{tr}_3 X_3\delta_i = \sum_{j=1}^{L}(X_3)_{ji} = 1 \tag{43}$$

we also get

$$\begin{aligned} (X_1 \otimes X_2)p'^{(2)} &= (X_1 \otimes X_2)\sum_{i=1}^{L} p'(i)\delta_i \otimes \delta_i && (44)\\ &= (X_1 \otimes X_2)\sum_{i=1}^{L} p'(i)\mathrm{tr}_3\delta_i^{\otimes 2} \otimes X_3\delta_i && (45)\\ &= \mathrm{tr}_3(X_1 \otimes X_2 \otimes X_3)\sum_{i=1}^{L} p'(i)\delta_i^{\otimes 3} && (46)\\ &= \mathrm{tr}_3(X_1 \otimes X_2 \otimes X_3)p'^{(3)}, && (47) \end{aligned}$$

thus we can use the assumed validity of the equation $p^{(3)} = (X_1 \otimes X_2 \otimes X_3)p'^{(3)}$

$$\begin{aligned} p^{(2)} &= \mathrm{tr}_3 p^{(3)} && (48)\\ &= \mathrm{tr}_3(X_1 \otimes X_2 \otimes X_3)p'^{(3)} && (49)\\ &= (X_1 \otimes X_2)p'^{(2)}, && (50) \end{aligned}$$

as desired. The same argument works with any of the other pairs. Our goal is to show first that the validity of these 3 pairwise equations already ensures that $X_1 = X_2 = X_3$. We argue as follows. First, fix $i \in [3]$ and choose any $j \neq i$ in $[3]$. Then define the matrix $\tilde{X}_i$ by $(\tilde{X}_i)_{mn} := p'(n)(X_i)_{mn}$. Note that the columns of $\tilde{X}_i$ form a linearly independent set (since X_i is invertible), and thus $\tilde{X}_i$ is invertible as well. For this argument to be valid, we do of course require our assumption $p(i) > 0\ \forall i \in [L]$ to hold true. We may now rewrite above equation slightly:

$$\begin{aligned} p^{(2)} &= X_i \otimes X_j(p'^{(2)}) && (51)\\ &= (\mathbb{1} \otimes X_j)(X_i \otimes \mathbb{1})(p'^{(2)}) && (52)\\ &= (\mathbb{1} \otimes X_j)\sum_{m=1}^{L} p'(m)(X_i \otimes \mathbb{1})\delta_m^{\otimes 2} && (53)\\ &= (\mathbb{1} \otimes X_j)\sum_{m,n=1}^{L} p'(m)(X_i)_{nm}\delta_n \otimes \delta_m && (54)\\ &= (\mathbb{1} \otimes X_j)\sum_{m,n=1}^{L} (\tilde{X}_i)_{nm}\delta_n \otimes \delta_m && (55)\\ &= (\mathbb{1} \otimes X_j)\sum_{n=1}^{L} \delta_n \otimes \tilde{X}_i^{\top}\delta_n && (56)\\ &= (\mathbb{1} \otimes X_j)(\mathbb{1} \otimes \tilde{X}_i^{\top})\sum_{n=1}^{L} \delta_n^{\otimes 2} && (57)\\ &= (\mathbb{1} \otimes (X_j \circ \tilde{X}_i^{\top}))\sum_{n=1}^{L} \delta_n^{\otimes 2}. && (58) \end{aligned}$$

It is clear that $X_j \circ \tilde{X}_i^{\top}$ is still invertible, hence it seems convenient for the moment to have a look at the equation

$$\sum_{i=1}^{L} p(i)\delta_i^{\otimes 2} = (\mathbb{1} \otimes X)\sum_{i=1}^{L} \delta_i^{\otimes 2}, \qquad X \in Gl(L). \tag{59}$$

Writing this out in coordinates immediately yields

$$\begin{aligned} p(i) &= \sum_{m=1}^{L} \delta(i,m)X_{im} = X_{ii} \qquad \forall i \in [L] && (60)\\ 0 &= \sum_{m=1}^{L} \delta(i,m)X_{jm} = X_{ij} \qquad \forall i \neq j \in [L]. && (61) \end{aligned}$$

Therefore, X is uniquely determined through p in that equation. It follows that all the $X_j \circ \tilde{X}_i^{\top} = Z_i$ for some Z_i, independent from the choice of j, and hence $X_j = Z_i \circ (\tilde{X}_i^{\top})^{-1}$ (independent from the choice of j). Playing this trick two times shows that, in fact, $X_1 = X_2 = X_3 =: X$ for some $X \in Gl_1(L)$.

We can now proceed to the second part of our proof.

Consider the equations

$$p^{(t)} = X^{\otimes t}p'^{(t)}, \qquad t \in [3], \tag{62}$$

where $p, p' \in \mathfrak{P}([L])$ and $X \in Gl_1(L)$. What are possible solutions of equation (62) if we require $p \in \mathfrak{P}_{>}([L])$ and $X \in Gl_1([L])$?

First, let us treat p and p' as fixed, and ask for solutions X. Define matrices $A, B \in Gl(L)$ by their entries $a_{ij} := \delta(i,j)\sqrt{p(i)}$ and $b_{ij} := \delta(i,j)\sqrt{p'(i)}$. Then we can set $t = 2$ and reformulate our equation (62) as

$$\sum_{i=1}^{L} \delta_i^{\otimes 2} = (A^{-1}XB \otimes A^{-1}XB)\sum_{i=1}^{L} \delta_i^{\otimes 2}. \tag{63}$$

Thus a more convenient formulation of problem (62) is to solve the equation

$$\sum_{i=1}^{L} \delta_i^{\otimes 2} = (Y \otimes Y) \sum_{i=1}^{L} \delta_i^{\otimes 2} \tag{64}$$

for Y, and we can now use the same trick as before and obtain $YY^\top = \mathbb{1}$. But this is equivalent to stating that Y is orthogonal! Rewinding things, we know now that there exists an orthogonal matrix Y such that $X = AYB^{-1}$.
We now use equation (62) with $t = 3$. We then get, using the special form of X that we just obtained, the equation

$$\sum_{i=1}^{L} \frac{1}{\sqrt{p(i)}} \delta_i^{\otimes 3} = Y^{\otimes 3} \sum_{i=1}^{L} \frac{1}{\sqrt{p'(i)}} \delta_i^{\otimes 3}. \tag{65}$$

Looking at specific entries, we see that the following are valid.

$$\frac{1}{\sqrt{p(i)}} = \sum_{j=1}^{L} \frac{1}{\sqrt{p'(j)}} Y_{ij}^2 Y_{ij} \qquad \forall i \in [L] \tag{66}$$

$$0 = \sum_{j=1}^{L} \frac{1}{\sqrt{p'(j)}} Y_{ij}^2 Y_{kj} \qquad \forall i \neq k \in [L]. \tag{67}$$

However, the vectors $v_i := \sum_{j=1}^{L} Y_{ij} \delta_j$ ($i \in [L]$) form an orthonormal set (since Y is orthogonal). Define the vectors $w_i := \sum_{j=1}^{L} \frac{1}{\sqrt{p'(j)}} Y_{ij}^2 \delta_j$ ($i \in [L]$), then the above can be reformulated as

$$\frac{1}{\sqrt{p(i)}} = \langle w_i, v_i \rangle \qquad \forall i \in [L] \tag{68}$$

$$0 = \langle w_i, v_k \rangle \qquad \forall i \neq k \in [L]. \tag{69}$$

But this clearly implies that, for each $i \in [L]$, we have the equalities

$$\frac{\sqrt{p'(j)}}{\sqrt{p(i)}} Y_{ij} = Y_{ij}^2 \qquad \forall j \in [L]. \tag{70}$$

Whenever a Y_{ij} is zero, these are trivially true. If $Y_{ij} \neq 0$, then we may divide by it and therefore obtain that, for each $i, j \in [L]$, either

$$Y_{ij} = 0 \qquad \text{or else} \qquad Y_{ij} = \sqrt{\frac{p'(j)}{p(i)}}. \tag{71}$$

Now we introduce sets $\mathcal{I}_i \subset [L] := \{j : Y_{ij} \neq 0\},\ i \in [L]$ such that the vectors v_i fulfill

$$v_i = \sum_{j \in I_i} \sqrt{\frac{p'(j)}{p(i)}} \delta_i. \tag{72}$$

Of course, since Y is an orthonormal matrix, these vectors again form an orthonormal set. Also, it holds that each $\mathcal{I}_i$ satisfies $|\mathcal{I}_i| > 0$. Hence for $i \neq l$,

$$0 = \sum_{j \in \mathcal{I}_i \cap \mathcal{I}_l} \sqrt{\frac{p'(j)}{p(i)}} \sqrt{\frac{p'(j)}{p(l)}} = \sum_{j \in \mathcal{I}_i \cap \mathcal{I}_l} p(j). \tag{73}$$

It follows that $\mathcal{I}_j \cap \mathcal{I}_k = \emptyset$, whenever $j \neq k$. Clearly then, since each of the sets $\mathcal{I}_i$ is non-empty, they all must be one-element sets. This also directly implies that, for $i \in \mathcal{I}_j$, it holds $Y_{ij} = 1$ (meaning that Y is a permutation matrix) and, additionally, that $p(i) = p'(j)$ whenever $i \in \mathcal{I}_j$. The desired permutation is conveniently defined through its action on $\mathfrak{P}([L])$:

$$\tau(\delta_i) := \mathbb{1}_{\mathcal{I}_i}. \tag{74}$$

This proves the equation (27) in Theorem 1.
How validity of (27) implies the existence of an inverse to the map defined in (28) can be seen as follows: Assume there are choices $W_1 \ldots, W_K \in \mathcal{W}([L], [L])$ and $V_1, \ldots, V_K \in \mathcal{W}([L], [L])$ of invertible channels and two distributions $p, p' \in \mathfrak{P}_>([L])$. Then our original question that we formalized in (6) can be recast as follows: Can it happen that the equality

$$\left(\bigotimes_{i=1}^{K} V_i \right) \sum_{i=1}^{L} p(i) \delta_i^{\otimes K} = \left(\bigotimes_{i=1}^{K} W_i \right) \sum_{i=1}^{L} p'(i) \delta_i^{\otimes K} \tag{75}$$

holds? Since all channels are invertible, one can equivalently ask whether the equation

$$\left(\bigotimes_{i=1}^{K} V_i^{-1} \circ W_i \right) \sum_{i=1}^{L} p'(i) \delta_i^{\otimes K} = \sum_{i=1}^{L} p(i) \delta_i^{\otimes K} \tag{76}$$

holds. The answer to this question has already been given now and lets us conclude that there is a permutation $\tau \in S_L$ such that $V_i^{-1} \circ W_i = \tau^{-1}$ for all $i = 1, \ldots, K$ and $p' = \tau(p)$. In particular, this implies that for all $i \in [K]$ we have

$$V_i = W_i \circ \tau, \tag{77}$$

so that the channels differ only by one joint permutation. If one additionally assumes that $p, p' \in \mathfrak{P}^{\downarrow}([L]) \cap \mathfrak{P}_>([L])$, then it follows that (76) can only hold if τ is the identity map on $[L]$ and, therefore, it holds that $V_i = W_i$ for all $i \in [K]$ and $p = p'$. Thus, the function Θ is injective.
It remains to define Θ': Define, for $q \in \mathfrak{P}([L]^K)$, $\Theta'(p)$ to be the choice $p \in \mathfrak{P}_>([L])$ that achieves the minimum in

$$\min_{p \in \mathfrak{P}_>([L])} \min_{W_1, \ldots, W_K} \|(\otimes_{i=1}^{K} W_i) p^{(K)} - q\|_1. \tag{78}$$

It is understood that the minimization is over invertible channels $W_1, \ldots, W_K$ only. Under the same assumptions on $p, W_1, \ldots, W_K$ as those made for the definition of Θ', the output $\Theta''(q)$ of the extended version Θ'' is conveniently defined as the minimizer $(p, W_1, \ldots, W_K)$ in

$$\min_{(p, W_1, \ldots, W_K)} \|(\otimes_{i=1}^{K} W_i) p^{(K)} - q\|_1. \tag{79}$$

One may replace $\| \cdot - q\|_1$ by any other norm in the above definitions, or by Kullback-Leibler divergence $D(\cdot||q)$, if additional requirements in the computation of Θ' or Θ'' are required. ∎

We will now deliver the proofs for situations where the output systems are strictly smaller than the input. The basic idea is that the observing agents report events that are elements of a system with output $L' \leq L$. The question then is, whether their findings are compatible with the assumption that $L = L'$.

Proof of Theorem 2: Let the output distribution of the system be $q \in \mathfrak{P}([L'])^K$, and assume that $L' \leq L$ and that there are invertible channels $W_1, \ldots, W_K \in \mathcal{W}([L'], [L'])$ and

$r \in \mathfrak{P}([L'])$ and $V_1, \ldots, V_K \in \mathcal{W}([L],[L'])$ each having full range ($\dim \mathrm{ran}(W_i) = \mathbb{R}^{L'}$) and $s \in \mathfrak{P}([L])$ such that both

$$q = (\otimes_{i=1}^K W_i) r^{(K)} \quad \text{and} \quad q = (\otimes_{i=1}^K V_i) s^{(K)}. \tag{80}$$

Then, define $R := \sum_{i=1}^{L'} r(i)^{1/2} E_{ii}$ and $S := \sum_{i=1}^{L} s(i)^{1/2} E_{ii}$ and $X_i := R^{-1} \circ W_i^{-1} \circ V_i \circ S$. Clearly every X_i is an element of $M(L, L')$. It holds, by playing the same tricks as before:

$$X_i X_j^\top = \mathbb{1}_{\mathbb{C}^{L'}}, \qquad \forall\ i \neq j \in [K]. \tag{81}$$

Now assume $K \geq 3$, and pick the indices $1, 2, 3$ as an example. Let the columns of $X_3^\top$ be $c_1, \ldots, c_{L'}$. These are linearly independent, since X_3 is full-range by assumption. They span the subspace $C \subset \mathbb{R}^L$ with $\dim C = L'$. From above equations (81) they further fulfill

$$X_i c_j = e_j, \qquad i = 1, 2,\ j = 1, \ldots, L'. \tag{82}$$

But that implies that $X_1 c = X_2 c$ for all $c \in C$. Since $\dim \mathrm{ran}\, X_i = L'$ we get $X_1 = X_2$, hence especially $\ker X_1 = \ker X_2$. This directly implies $\ker V_1 = \ker V_2$. The same argument holds with any different choice of indices $i, j, k \in [K]$, thus the desired

$$\ker V_1 = \ldots = \ker V_K \tag{83}$$

follows. ∎

Proof of Theorem 3: If $L' > L$, and $K \geq 3$, we can use all our previous reasoning to do the following steps: First, for two $DCA(L', K, L)$ systems $(p, W_1, \ldots, W_K)$ and $(p^\backslash, V_1, \ldots, V_K)$ we get

$$W_i^{-1} \circ V_i = W_j^{-1} \circ V_j \qquad \forall i, j \in [K]. \tag{84}$$

With $X := \sqrt{p^{-1}} \circ W_1^{-1} \circ V_1 \circ \sqrt{p^\backslash}$ we then get, as before, that $X = \tau$ for some permutation τ on $[L']$, and it follows that the two systems are equivalent up to permutations. ∎

Proof of Theorem 4: We will first and in greater depth consider the case $L' = 2$. In a second step, we generalize our proof to arbitrary L'.
Let $\{\gamma_i\}_{i=1}^L \subset \mathbb{R}$ be such that

$$0 = \sum_{i=1}^L \gamma_i W^{\otimes K}(\delta_i^{\otimes K}). \tag{85}$$

We can assume that we have $W(\delta_i) = \lambda_i \delta_1 + \lambda_i' \delta_2$ for an appropriate set of $\{\lambda_i\}_{i=1}^L \subset [0,1]$ and with the convention $\lambda_i' := 1 - \lambda_i$ for all $i \in [L]$. We also know that the set $\{\delta_1^{\otimes t} \otimes \delta_2^{\otimes K-t}\}_{t=0}^K$ is linearly independent. Thus, we get for every $t \in \{0, \ldots, K\}$:

$$0 = \sum_{i=1}^L \gamma_i \lambda_i^t (\lambda_i')^{K-t}. \tag{86}$$

This set of equations can easily be seen to be equivalent to

$$0 = \sum_{i=1}^K \gamma_i \frac{1}{\binom{K}{t}} B_{t,K}(\lambda_i) \qquad \forall t = 0, \ldots, K, \tag{87}$$

where $B_{t,K}$ are the Bernstein-polynomials [1] defined as $B_{t,K}(x) := \binom{K}{t} x^t (1-x)^{K-t}$ ($x \in \mathbb{R}$). It is known[1] [15] that these polynomials span the space $Pl(K)$ of all polynomials of degree no more than K. We can therefore reformulate (87) as

$$0 = \sum_{i=1}^L \gamma_i P(\lambda_i) \qquad \forall P \in Pl(K). \tag{88}$$

Given a choice of the γ_i, we can always find a polynomial P satisfying $P(\lambda_i) = \gamma_i$, as long as $K - 1 \geq L$. Then, we are left with the equation

$$0 = \sum_{i=1}^L \gamma_i^2, \tag{89}$$

which clearly implies $\gamma_1 = \ldots = \gamma_L = 0$, hence $W^{\otimes K}$ is invertible.

In the more general setting where $L > L'$ and with W having full-range image we know there exist points $\{\mathbf{a}_i\}_{i=1}^L \subset \mathbb{R}^{L'-1}$ such that for the output distributions $W(\delta_i)$ it holds, with $a_{i,0} := 1 - \sum_{j=1}^{L'-1} a_{i,j}$: $W(\delta_i) = \sum_{j=1}^{L'-1} a_{i,j}\delta_j + a_{i,0}\delta_{L'}$.
Then as before, we can rewrite the statement

$$0 = \sum_{i=1}^L \gamma_i W^{\otimes K}(\delta_i^{\otimes K}) \tag{90}$$

as

$$0 = \sum_{i=1}^L \gamma_i B_{f,K}(\mathbf{a}_i) \qquad \forall\ f, \tag{91}$$

where $B_{f,K}$ are the multivariate Bernstein polynomials and $f : [L] \to \mathbb{N}$ are so-called frequencies and satisfy $f \geq 0$ as well as $\sum_{i=1}^L f(i) = K$. But for a fixed choice of the γ_i, this means that the statement

$$0 = \sum_i \gamma_i P(\mathbf{a}_i) \tag{92}$$

has to hold *for all* linear combinations P of the polynomials $B_{f,K}$. It is known [21] that these span the space of all polynomials of degree less than or equal to K. Hence, above equality carries over to all polynomials P in $L-1$ variables of degree less than or equal to K. The theory of Kergin-interpolation [22] (and [28]) tells us that we can always find a polynomial P_γ satisfying $P_\gamma(\mathbf{a}_i) = \gamma_i$, $i = 1, \ldots, L$, if $K \geq L - 1$ holds. ∎

We will now come to the last remaining one of our proofs.

Proof of Lemma 1: Just let each $W_1, \ldots, W_K = W$ where W is defined by

$$W(\delta_1) = \delta_1, \quad W(\delta_2) = \delta_2, \quad W(\delta_3) = W(\delta_4) = r, \tag{93}$$

where $r \in \mathfrak{P}([L])$ is arbitrary. Then, it is impossible to tell the values $p(4)$ and $p(3)$ of the sought-after p: think of the sets

$$P_{ab} := \{p : p(1) = a,\ p(2) = b\}. \tag{94}$$

[1]This fact was first brought to our attention through the "On-Line Geometric Modeling Notes" of Kenneth Joy at UC Davis.

Every of these distributions gets mapped to the same distribution

$$a \cdot \delta_1^{\otimes K} + b \cdot \delta_2^{\otimes K} + (1-a-b) r^{\otimes K} \otimes r^{\otimes K} \qquad (95)$$

by application of above channels. Thus, we can *never* hope to get an invertibility criterion just from looking at the pairwise mutual informations! ∎

VI. Appendix

We now briefly touch upon the topics hypothesis testing and statistical inference.

A. Hypothesis testing for dependent component systems

In this section, we present some facts on hypothesis testing in order to connect this presentation to it. Given an arbitrary DCS $\mathcal{S}$ with parameters L, K, L', what we receive at the output during $n \in \mathbb{N}$ observations is a string $y^n \in ([L']^n)$. Let the number of times each symbol $(y_1, \dots, y_K) \in [L']^K$ appears in $y^n = ((y_{1,1}, \dots, y_{K,1}), \dots, (y_{1,n}, \dots, y_{K,n}))$ be $N(y_1, \dots, y_K | y^n)$. Due to the memoryless nature of the system, every permutation of that string is equally likely to be the output of the system. Hence, the whole set $\{\tau y^n : \tau \in S_n\}$ (where S_n is the set of permutations on $n \in \mathbb{N}$ symbols and $(\tau y^n)_i := (y_{1,\tau^{-1}(i)}, \dots, y_{K,\tau^{-1}(i)})$) will get mapped to the same estimate $\hat{q} = \hat{q}(y^n)$. We will henceforth, for nonnegative functions $f : [L']^K \to \mathbb{N}$, use the abbreviation

$$T_f := \{y^n : N(\cdot|y^n) = f\}. \qquad (96)$$

Note that $T_{N(\cdot|y^n)} = \{\tau y^n : \tau \in S_n\}$ holds. What is the best (asymptotic) estimate, given y^n? This question is answered as follows: We search for

$$\hat{q} := \arg\max\{q^{\otimes n}(y^n) : q \in \mathfrak{P}([L'])^K)\}. \qquad (97)$$

According to [11, Proof of Lemma 2.3], the solution to this optimization problem is given by the maximum-likelihood estimate $\hat{q} = \frac{1}{n} N(\cdot|y^n)$. Moreover, the probability that this estimate $\hat{q}$ satisfies $D(\hat{q}\|q) > \varepsilon$ for some $\varepsilon > 0$ is bounded by

$$q^{\otimes n}(\{y^n : D(\tfrac{1}{n} N(\cdot|y^n)\|q) > \varepsilon\}) \leq poly(n) 2^{-n \cdot \varepsilon}, \qquad (98)$$

where $poly(n)$ denotes a polynomial in n that depends on L as well. Obviously, this upper bound goes to zero exponentially fast in n. One may choose ε depending on n by setting e.g. $\varepsilon_n := 1/\sqrt{n}$, and this delivers a hypothesis test that succeeds with probability going to one as n tends to infinity.

Adopted to our scenario, it will identify the output distribution q of any DCS with arbitrary precision. It remains to be proven that small errors in the estimate remain bounded when inverting the system. Also, the question of optimality of this choice of test remains.

In hypothesis testing scenarios as described for example in [4], one typically considers binary hypotheses. That is, one assumes that the true state of the system is given by either of two distribution r, s. In the scenario described in this work however, we are faced with an additional subtlety: The channels that map the system outputs to the observer are unknown as well.

Thus a rigorous problem formulation in our scenario includes the adoption of additional hypotheses on the channels, may these be of the nature 'it is either $\otimes_i W_i$ or $\otimes_i V_i$' or 'the channels are drawn at random according to some distribution on the set of channels'. A worst-case assumption would be that every of the channels $\otimes_i W_i$ is possible. In that case we encounter an additional problem: The quantity

$$\inf_{\mathfrak{W}, \mathfrak{V} \in \mathcal{W}_{\leftrightarrow}(L,K)} D(\mathfrak{W} r^{(K)} \| \mathfrak{V} s^{(K)}), \qquad (99)$$

where $\mathcal{W}_{\leftrightarrow}(L, K) \subset \mathcal{W}([L], [L])^{\otimes K}$ is the subset of invertible channels of the form $\mathfrak{W} = W_1 \otimes \dots \otimes W_K$ is simply equal to zero. This follows from the fact that in every vicinity of a non-invertible channel there is an invertible channel, too.

B. Connection to statistical inference

In this subsection we briefly connect our findings to the area of statistical inference. Let for simplicity $\mathbf{A}, \mathbf{B}, \mathbf{C}$ be binary alphabets. We let $\mathbf{C}$ be the input of a DCS and $\mathbf{A}$, $\mathbf{B}$ the output systems. Let $q \in \mathfrak{P}(\mathbf{A}, \mathbf{B})$ be the output of the DCS. Let the overall distribution be $s \in \mathfrak{P}(\mathbf{A} \times \mathbf{B} \times \mathbf{C})$. Due to the special structure of our system we have

$$s_{A|BC} = s_{A|C}. \qquad (100)$$

It is also clear by construction that the events happening on $\mathbf{C}$ are a common cause for those on $\mathbf{A}$ and $\mathbf{B}$. As explained in [16], [27] such systems go under the name 'conjunctive fork'. They are to be distinguished from systems which are included under the name 'Simpson's paradox'. While in our case we have two channels going (strictly speaking) in parallel from $\mathbf{C} \times \mathbf{C}$ to $\mathbf{A} \times \mathbf{B}$, a system on which Simpson's paradox (wrongly inferring that some statistical event is causal for another statistical event) can occur would for example have to have a Markov structure $\mathbf{C} \to \mathbf{A} \to \mathbf{B}$.

VII. Open problems

Stability: Once we have estimated the output distribution on $[L']^K$ we would like to invert it in order to know p. Then, if a small mistake in the estimation scheme would lead to dramatically different results for p, we would rightfully see this as a drastic drawback of the method. This issue deserves further attention.

Hypothesis testing: It would be desirable to compare different hypothesis testing scenarios and derive optimal tests for them. It remains to be seen whether this can lead to the derivation of new and meaningful information measures. For more information, see the appendix.

Activation: We left open the question of a general 'activation' effect of invertibility. It would be interesting to know under what conditions a general set of K channels $W_1, \dots, W_K \in \mathcal{W}([L], [L'])$ becomes invertible as a map $\otimes_{i=1}^K W_i$ from $\mathfrak{P}^{(K)}([L])$ to $\mathfrak{P}([L']^K)$, for general L and L'. This includes a more detailed study of the projective behaviour that was only touched upon in Remark 3.

Multivariate polynomials: A detailed investigation of this connection is postponed to future work.

Another possible route for future research is the connection of our findings to the very structure of three-dimensional space as experienced by us every day, as has been done e.g. in [25] for qubits.

At last we have to mention that, of course, an analysis of infinite-dimensional DCSs, extensions to quantum mechanics and an investigation of DCS with time varying channels or distributions offer the potential of finding results that are interesting in their own right.

REFERENCES

[1] S. N. Bernstein, "Démonstration du théorème de Weierstrass fondée sur le calcul des probabilités", Communications de la Société Mathématique de Kharkov, 2. Series XIII No. 1, 12. (1912)

[2] H. H. Beverage and H. O. Peterson, "Diversity receiving system of RCA communications, inc., for radiotelegraphy", *Proc. IRE*, vVol. 19, 531-561 (1931)

[3] H. O. Peterson, H. H. Beverage, J. B. Moore, "Diversity telephone receiving system of RCA communications, inc.", *Proc. IRE*, Vol. 19, 562-584 (1931)

[4] R.E. Blahut, "Hypothesis Testing and Information Theroy", *IEEE Trans. Inf. Theory*, Vol. IT-20 405-417 (1974)

[5] R. S. Blum, S. A. Kassam, H. V. Poor, "Distributed Detection with Multiple Sensors: Part II–Advanced Topics", *Proceedings of the IEEE*, Vol.85, No.1, 6-23 (1997)

[6] D. G. Brennan, "Linear diversity combining techniques", *Proc. IRE* Vol. 47, 1075-1102 (1959), Reprint: *Proc. IEEE* Vol. 91, No. 2, 331-356 (2003)

[7] P. Busch, "Informationally complete sets of physical quantities", *International Journal of Theoretical Physics*, Vol. 30, No. 9, 1217–1227 (1991)

[8] P. Comon, "Independent component analysis, A new concept?", *Signal Processing*, Vol. 36, Iss. 3, 287-314 (1994)

[9] H. Chernoff, "A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations", *Ann. Math. Statist.*, Vol. 23, 495-507 (1952)

[10] I. Csiszar, G. Longo, "On the error exponent for source coding and for testing simple statistical hypotheses", *Studia Sci. Math. Hungar.* Vol. 6, 181-191 (1971)

[11] I. Csiszar, J. Körner, *Information Theory; Coding Theorems for Discrete Memoryless Systems*, Akadémiai Kiadó, Budapest/Academic Press Inc., New York 1981

[12] G.M. D'Ariano, P. Perinotti, M.F. Sacchi, "Informationally complete measurements and groups representation", *J. Opt. B: Quantum Semiclass. Opt.* Vol. 6, 487-491 (2004)

[13] G. M. D'Ariano, P. Perinotti, M. F. Sacchi, "Quantum universal detectors", *EPL* Vol. 65, No. 2, 165 (2004)

[14] Y. Ephraim, N. Merhav, "Hidden Markov Models", *Trans. Inf. Theory*, Vol. 48, No. 6 (2002)

[15] R.T. Farouki, "The Bernstein polynomial basis: A centennial retrospective", *Computer Aided Geometric Design*, Vol. 29 Iss. 6, 379-419 (2012)

[16] B. V. Frosini, "Causality and Causal Models: A Conceptual Perspective", *International Statistical Review / Revue Internationale de Statistique*, Vol. 74, No. 3, 305-334 (2006)

[17] G.M. Gemelos, S. Sigurjónsson, T. Weissmann, "Universal Minimax Discrete Denoising Under Channel Uncertainty", *Trans. Inf. Theory*, Vol. 52, No. 8 (2006)

[18] T.S. Han, S. Verdú, "Approximation Theory of Output Statistics", *Trans. Inf. Theory*, Vol. 39, No. 3 (1993)

[19] T.S. Han, S. Verdú, "Spectrum Invariancy under Output Approximation for Discrete Memoryless Channels with Full Rank", *Problemii Peredachi Informatsii*, (in Russian), Vol. 29, No. 2, 9-27 (1993) translated in *Problems of Information Transmission*, 101-118 (1993)

[20] A. Hyvärinen, E. Oja,"Independent component analysis: algorithms and applications", *Neural Networks*, vol. 13, 411-430 (2000)

[21] K. Jetter, J. Stöckler, "An identity for multivariate Bernstein polynomials", *Computer Aided Geometric Design*, Vol. 20 Iss. 8-9, 563-577 (2003)

[22] P. Kergin, "A natural interpolation of functions" *J. Approx. Th.* Vol. 29, 278293 (1980)

[23] S. Kullback, "Information Theory and Statistics", *New York: Dover* (1968) and *New York: Wiley* (1959)

[24] R. Li, H. Li, "Dependent Component Analysis: Concepts and Main Algorithms", *Journal of Computers* Vol. 5, No. 4 (2010)

[25] M.P. Müller, L. Masanes, "Three-dimensionality of space and the quantum bit: an information-theoretic approach", *New J. Phys.*, Vol. 15, 053040 (2013)

[26] E. Prugovečki, "Information-theoretical aspects of quantum measurement", *International Journal of Theoretical Physics*, Vol. 16, No. 5, 321–331 (1977)

[27] H. Reichenbach, "The Direction of Time", *Berkeley: University of California Press* (1956)

[28] T. Sauer, Y. Xu, "On Multivariate Lagrange Interpolation" *Math. of Comp.* Vol. 64, No. 211, 1147-1170 (1995)

[29] R. Viswanathan, P. K. Varshney, "Distributed Detection With Multiple Sensors: Part I–Fundamentals", *Proceedings of the IEEE*, Vol.85, No.1, 54-63 (1997)

[30] C.F. von Weizsäcker, "The Structure of Physics", *Springer Verlag, Dordrecht* (2006)

[31] T. Weissmann, E. Ordentlich, G. Seroussi, S. Verdú, M.J. Weinberger, "Universal Discrete Denoising: Known Channel", *Trans. Inf. Theory*, Vol. 51, No. 1 (2005)

[32] S. YiChang, V. Kwatra, T. Chinen, F. Hui, S. Ioffe, "Joint Noise Level Estimation from Personal Photo Collections", *Computer Vision (ICCV), 2013 IEEE International Conference on*, 2896-2903 (2013)

[33] Y. Xiang, D. Peng, Z. Yang, *Blind Source Separation - Dependent Component Analysis*, SpringerBriefs in Electrical and Computer Engineering (2015)

Janis Nötzel received the Dipl. Phys. degree in physics from the Technische Universität Berlin, Germany, in 2007 and the PhD degree from Technische Universität München, Germany, in 2012. From 2008 to 2010 he was a research assistant at the Technische Universität Berlin, Germany, and from 2011 until 2015 at Technische Universität München. From July 2015 to August 2016 he was a DFG research fellow at Universitat Autònoma de Barcelona, Spain. He is now with the Technische Universität Dresden, Communications Laboratory, 01069 Dresden, Germany.

Walter Swetly studied Electrical Engineering, Logic, Philosophy and Theoretical Linguistics at Technische Universität Mnchen and Ludwig-Maximilians-Universität Mn̈chen. He received his PhD degree from LMU Mn̈chen in Logic in 2010. From 2010 - 2012 he was a member of the Chair for Logic and Philosophy of Science at LMU München and from 2012 - 2014 at the Chair for Data Processing at Technische Universität München. Since then he has worked in various positions in the industry.